

A Time Consistent Formulation of Risk Constrained Stochastic Optimal Control

Yin-Lam Chow, Marco Pavone

Abstract—Time-consistency is an essential requirement in risk sensitive optimal control problems to make rational decisions. An optimization problem is time consistent if its solution policy does not depend on the time sequence of solving the optimization problem. On the other hand, a dynamic risk measure is time consistent if a certain outcome is considered less risky in the future implies this outcome is also less risky at current stage. In this paper, we study time-consistency of risk constrained problem where the risk metric is time consistent. From the Bellman optimality condition in [1], we establish an analytical “risk-to-go” that results in a time consistent optimal policy. Finally we demonstrate the effectiveness of the analytical solution by solving Haviv’s counter-example [2] in time inconsistent planning.

I. INTRODUCTION

Stochastic Optimal Control (SOC) is concerned with sequential decision-making under uncertainty. Consider a dynamical process that can be influenced by exogenous noises as well as decisions made at every time step. The decision maker wants to optimize the behavior of the dynamical system over a certain time horizon by finding a policy that maps the history of states to optimal actions.

In this SOC setup, we solve an optimization problem at present time and determine an optimal policy that maps states to actions subsequently to minimize cumulative cost. In order to obtain rational decisions, we aim to solve for a *time consistent control policy* for which this solution policy is optimal to both the SOC at present and the tail-subproblems in all subsequent time steps. If such policy exists, the SOC problem is also known to be time consistent. The property of time consistent problem is formally stated as follows. The decision maker formulates an optimization problem at time $k = 0$ that yields a sequence of optimal decision rules for $k = 0$ to $k = N$. Then, at the next time step $k = 1$, he/she formulates a new problem starting at $k = 1$ that yields a new sequence of optimal decision rules from time steps $k = 1$ to N . The sequence of policy is *time consistent* if the strategies obtained when solving the original problem at time $k = 0$ remain optimal for all subsequent problems.

Recently, the concept of time-consistency has also been extended to the context of risk measures [3], [4], [5], [6], [7], [8], [9]. In these papers, the authors formally defined the notion of time-consistency of risks, provided examples of time consistent risk measures and axiomatically justified that this property is necessary to develop rational risk assessments in stochastic processes. In [10], the author showed that expectation and worst case risk are the only time consistent coherent risks, and the authors in [11] developed a necessary and sufficient condition for time consistent risk measures. Furthermore in [12], the authors provided tight approxima-

tions of time inconsistent coherent risks by lower or upper-bounding them with time consistent metrics.

Note that common examples of time consistent risk measures include expectation and entropic risk measures. From [11], by posing an unconstrained stochastic optimal control problem with a time consistent risk measure, one obtains a time consistent solution policy by dynamic programming. However [13] shows that a risk constrained problem is not necessarily time consistent even if both objective function and constraints are time consistent risk measures. This results in undesirable outcomes for example when a decision maker seeks to minimize expected loss subjected to a risk constraint, the optimal policy at present may become infeasible in future when the risk is re-evaluated.

While time consistent policies are essential to ensure rational decisions, it has been pointed out in [2], [14], [15] that in a constrained SOC setup, time consistency is not necessarily satisfied by an optimal policy. There are several sufficient conditions to guarantee time consistency for specific constrained SOC problems. In [16], the author provided a sufficient condition for time consistency in deterministic optimal control problems. Also, [13] showed that the risk constrained problem is time consistent if the risk measures are *optimality consistency*, i.e., any constraints that are feasible at present will also be also feasible in future. Furthermore [2], [14] argued that a SOC problem is time consistent if constraints are satisfied at every sample history path. However the above sufficient conditions are either restricted to a small SOC problem subclass or verifying this condition requires exponential computational complexity. In contrary the results in this paper shed light to a simple and analytic sufficient condition of time-consistency for a general class of risk constrained SOC problems.

Our contributions of this paper are three-fold.

- First, we formulate a risk constrained SOC problem with time consistent risk measures and show that it can be solved by dynamic programming techniques (in the augmented action space).
- Second, by reformulating the above problem into an augmented Markov decision problem, we develop an analytical solution for the “risk-to-go” that results a time consistent optimal control policy.
- Third, we illustrate the effectiveness of this analytical method by solving for a time consistent optimal policy to Haviv’s “squander or save” counter-example [2] on time-inconsistent planning.

The rest of the paper is organized as follows. In Section II we provide a review of the theory of risk metrics and Markov decision processes. In Section III we provide an analytical solution to the risk-to-go update that yields a time consistent optimal control policy. In Section IV we further justify our solution method by solving for a time consistent

optimal policy to the “squander or save” problem. Finally, the conclusion and future work are discussed in Section V.

II. PRELIMINARIES

In this section we provide some background for dynamic, time-consistent risk metrics and risk constrained SOC problems, on which we will rely extensively later in the paper.

A. Markov Decision Processes

A finite Markov Decision Process (MDP) is a four-tuple $(S, U, Q, U(\cdot))$, where S , the state space, is a finite set; U , the control space, is a finite set; for every $x \in S$, $U(x) \subseteq U$ is a nonempty set which represents the set of admissible controls when the system state is x ; and, finally, $Q(\cdot|x, u)$ (the transition probability) is a conditional probability on S given the set of admissible state-control pairs, i.e., the sets of pairs (x, u) where $x \in S$ and $u \in U(x)$.

Define the space H_k of admissible histories up to time k by $H_k = H_{k-1} \times S \times U$, for $k \geq 1$, and $H_0 = S$. A generic element $h_{0,k} \in H_k$ is of the form $h_{0,k} = (x_0, u_0, \dots, x_{k-1}, u_{k-1}, x_k)$. Let Π be the set of all deterministic policies with the property that at each time k the control is a function of $h_{0,k}$. In other words, $\Pi := \{\{\pi_0 : H_0 \rightarrow U, \pi_1 : H_1 \rightarrow U, \dots\} | \pi_k(h_{0,k}) \in U(x_k) \text{ for all } h_{0,k} \in H_k, k \geq 0\}$.

B. Dynamic, Time-consistent, Risk Measures

Consider a probability space (Ω, \mathcal{F}, P) , a filtration $\mathcal{F}_1 \subset \mathcal{F}_2 \dots \subset \mathcal{F}_N \subset \mathcal{F}$, and an adapted sequence of random variables $Z_k, k \in \{0, \dots, N\}$. We assume that $\mathcal{F}_0 = \{\Omega, \emptyset\}$, i.e., Z_0 is deterministic. In this paper we interpret the variables Z_k as stage-wise costs. For each $k \in \{1, \dots, N\}$, define the space of random variables with finite p th order moment as $\mathcal{Z}_k := L_p(\Omega, \mathcal{F}_k, P)$, $p \in [1, \infty]$; also, let $\mathcal{Z}_{k,N} := \mathcal{Z}_k \times \dots \times \mathcal{Z}_N$.

A dynamic risk measure is a sequence of monotone mappings $\rho_{k,N} : \mathcal{Z}_{k,N} \rightarrow \mathcal{Z}_k, k \in \{0, \dots, N\}$. Roughly speaking, a dynamic risk measure is *time consistent* if it is such that, when a Z cost sequence is deemed less risky than a W cost sequence from the perspective of a future time k , and both sequences yield identical costs from the current time l to the future time k , then the Z sequence is deemed less risky at the current time l , as well. We refer to [11] for a formal definition of time consistency. It turns out that dynamic, time-consistent risk metrics can be constructed by “compounding” coherent one-step conditional risk measures, which are defined as follows.

Definition II.1 (Coherent One-step Conditional Risk Measures). A coherent one-step conditional risk measures is a mapping $\rho_k : \mathcal{Z}_{k+1} \rightarrow \mathcal{Z}_k, k \in \{0, \dots, N\}$, with the following four properties:

- *Convexity*: $\rho_k(\lambda Z + (1 - \lambda)W) \leq \lambda \rho_k(Z) + (1 - \lambda) \rho_k(W), \forall \lambda \in [0, 1]$ and $Z, W \in \mathcal{Z}_{k+1}$;
- *Monotonicity*: if $Z \leq W$ then $\rho_k(Z) \leq \rho_k(W), \forall Z, W \in \mathcal{Z}_{k+1}$;
- *Translation invariance*: $\rho_k(Z + W) = Z + \rho_k(W), \forall Z \in \mathcal{Z}_k$ and $W \in \mathcal{Z}_{k+1}$;
- *Positive homogeneity*: $\rho_k(\lambda Z) = \lambda \rho_k(Z), \forall Z \in \mathcal{Z}_{k+1}$ and $\lambda \geq 0$.

The compositional structure of dynamic, time-consistent risk metrics is then characterized by the following theorem.

Theorem II.2 (Dynamic, Time-consistent Risk Metrics [11]). Consider, for each $k \in \{0, \dots, N\}$, the mappings $\rho_{k,N} : \mathcal{Z}_{k,N} \rightarrow \mathcal{Z}_k$ defined as

$$\rho_{k,N} = Z_k + \rho_k(Z_{k+1} + \rho_{k+1}(Z_{k+2} + \dots + \rho_{N-2}(Z_{N-1} + \rho_{N-1}(Z_N)) \dots)), \quad (1)$$

where the ρ_k 's are coherent one-step conditional risk measures. Then, the ensemble of such mappings is a dynamic, time-consistent risk measure.

In this paper we consider a (slight) refinement of the concept of dynamic, time-consistent risk metric, which involves the addition of a Markovian structure [11] and enables the development of dynamic programming equations.

Definition II.3 (Markov Dynamic Risk Measures [11]). Let $\mathcal{V} := L_p(S, \mathcal{B}, P)$ be the space of random variables on S with finite p th moment. Given a controlled Markov process $\{x_k\}$, a dynamic, time-consistent risk metric is a Markov dynamic risk metric if each coherent one-step conditional risk measure $\rho_k : \mathcal{Z}_{k+1} \rightarrow \mathcal{Z}_k$ in (1) can be written as:

$$\rho_k(V(x_{k+1})) = \sigma_k(V(x_{k+1}), x_k, Q(x_{k+1}|x_k, u_k)), \quad (2)$$

for all $V(x_{k+1}) \in \mathcal{V}$ and $u \in U(x_k)$, where σ_k is a coherent one-step conditional risk measure on \mathcal{V} (with the additional technical property that for every $V(x_{k+1}) \in \mathcal{V}$ and $u \in U(x_k)$ the function $x_k \mapsto \sigma_k(V(x_{k+1}), x_k, Q(x_{k+1}|x_k, u_k))$ is an element of \mathcal{V}).

In other words, in a Markov dynamic risk measures, the evaluation of risk is *not* allowed to depend on the whole past.

C. Stochastic Optimal Control with Dynamic, Time-consistent Risk Constraints

Consider an MDP and let $c : S \times U \rightarrow \mathbb{R}$ and $d : S \times U \rightarrow \mathbb{R}$ be functions which denote costs associated with state-action pairs. Given a policy $\pi \in \Pi$, an initial state $x_0 \in S$, and an horizon $N \geq 1$, the multi-stage cost function is defined as

$$J_N^\pi(x_0) := \mathbb{E} \left[\sum_{k=0}^{N-1} c(x_k, u_k) \right],$$

and the risk constraint is defined as

$$R_N^\pi(x_0) := \rho_{0,N} \left(d(x_0, u_0), \dots, d(x_{N-1}, u_{N-1}), 0 \right),$$

where $\rho_{k,N}(\cdot), k \in \{0, \dots, N-1\}$, is a Markov dynamic risk metric (for simplicity, we do not consider terminal costs, even though their inclusion is straightforward). The problem is then as follows:

Optimization problem OPT — Given an initial state $x_0 \in S$, a time horizon $N \geq 1$, and a risk threshold $r_0 \in \mathbb{R}$, solve

$$\begin{aligned} \min_{\pi \in \Pi} \quad & J_N^\pi(x_0) \\ \text{subject to} \quad & R_N^\pi(x_0) \leq r_0. \end{aligned}$$

If problem OPT is not feasible, we say that its value is ∞ . In [1] the authors developed a dynamic programming approach to solve this problem. To define the value functions, one needs to define the tail subproblems. For a given $k \in \{0, \dots, N-1\}$ and a given state $x_k \in S$, we define the *sub-histories* as $h_{k,j} := (x_k, u_k, \dots, x_j)$ for $j \in \{k, \dots, N\}$; also, we define the *space of truncated policies*

as $\Pi_k := \left\{ \{\pi_k, \pi_{k+1}, \dots\} \mid \pi_j(h_{k,j}) \in U(x_j) \text{ for } j \geq k \right\}$. For a given stage k and state x_k , the cost of the tail process associated with a policy $\pi \in \Pi_k$ is simply $J_N^\pi(x_k) := \mathbb{E} \left[\sum_{j=k}^{N-1} c(x_j, u_j) \right]$. The risk associated with the tail process is:

$$R_N^\pi(x_k) := \rho_{k,N} \left(d(x_k, u_k), \dots, d(x_{N-1}, u_{N-1}), 0 \right).$$

The tail subproblems are then defined as

$$\min_{\pi \in \Pi_k} J_N^\pi(x_k) \quad (3)$$

$$\text{subject to } R_N^\pi(x_k) \leq r_k(x_k), \quad (4)$$

for a given (undetermined) threshold value $r_k(x_k) \in \mathbb{R}$ (i.e., the tail subproblems are specified up to a threshold value).

For each $k \in \{0, \dots, N-1\}$ and $x_k \in S$, we define the set of feasible constraint thresholds as

$$\Phi_k(x_k) := [\underline{R}_N(x_k), \infty), \quad \Phi_N(x_N) := [0, \infty),$$

where $\underline{R}_N(x_k) := \min_{\pi \in \Pi_k} R_N^\pi(x_k)$. One then defines the value functions as follows:

- If $k < N$ and $r_k \in \Phi_k(x_k)$:

$$V_k(x_k, r_k) = \min_{\pi \in \Pi_k} J_N^\pi(x_k) \quad \text{subject to } R_N^\pi(x_k) \leq r_k.$$

- If $k \leq N$ and $r_k \notin \Phi_k(x_k)$:

$$V_k(x_k, r_k) = \infty.$$

- When $k = N$ and $r_N \in \Phi_N(x_N) = [0, \infty]$:

$$V_N(x_N, r_N) = 0.$$

Let $B(S)$ denote the space of real-valued bounded functions on S , and $B(S \times \mathbb{R})$ denote the space of real-valued bounded functions on $S \times \mathbb{R}$. For $k \in \{0, \dots, N-1\}$, we define the dynamic programming operator $T_k[V_{k+1}] : B(S \times \mathbb{R}) \mapsto B(S \times \mathbb{R})$ according to the equation:

$$T_k[V_{k+1}](x_k, r_k) := \inf_{(u, r') \in F_k(x_k, r_k)} \left\{ c(x_k, u) + \sum_{x_{k+1} \in S} Q(x_{k+1} | x_k, u) V_{k+1}(x_{k+1}, r'(x_{k+1})) \right\}, \quad (5)$$

where F_k is the set of control/threshold functions:

$$F_k(x_k, r_k) := \left\{ (u, r') \mid u \in U(x_k), r'(x') \in \Phi_{k+1}(x') \text{ for all } x' \in S, \text{ and } d(x_k, u) + \rho_k(r'(x_{k+1})) \leq r_k \right\}.$$

If $F_k(x_k, r_k) = \emptyset$, then $T_k[V_{k+1}](x_k, r_k) = \infty$.

For a given state and threshold constraint, F_k characterizes the set of feasible pairs of actions and subsequent constraint thresholds. Feasible subsequent constraint thresholds are thresholds which if satisfied at the next stage ensure that the current state satisfies the given constraint threshold. Note that the value functions are defined on an *augmented* state space, which combines the original (discrete) states x_k with the *real-valued risk-to-go* states r_k . We will refer to the

MDP problem associated with such augmented state space as augmented MDP (AMDP). The main result in [1] is the following theorem about the correctness of value iteration for AMDP.

Theorem II.4 (Bellman's Equation with Risk Constraints [1]). *For all $k \in \{0, \dots, N-1\}$ the value functions satisfy the Bellman's equation:*

$$V_k(x_k, r_k) = T_k[V_{k+1}](x_k, r_k).$$

Next, we present a procedure to construct optimal policies. Under the assumptions of Theorem II.4, for any given $x_k \in S$ and $r_k \in \Phi_k(x_k)$ (which implies that $F_k(x_k, r_k)$ is non-empty), let $u^*(x_k, r_k)$ and $r'(x_k, r_k)(\cdot)$ be the minimizers in equation (5). Next theorem shows how to construct history dependent optimal policies.

Theorem II.5 (Optimal Policies). *Let $\pi \in \Pi$ be a policy recursively defined as:*

$$\pi_k(h_k) = u^*(x_k, r_k), \text{ with } r_k = r'(x_{k-1}, r_{k-1})(x_k),$$

when $k \in \{1, \dots, N-1\}$, and

$$\pi(x_0) = u^*(x_0, r_0),$$

for a given threshold $r_0 \in \Phi_0(x_0)$. Then, π is an optimal policy for problem \mathcal{OPT} with initial condition x_0 and constraint threshold r_0 .

Interestingly, if one views the constraint thresholds as state variables (whose dynamics are given in the statement of Theorem II.5), the optimal (history-dependant) policies of problem \mathcal{OPT} have a Markovian structure with respect to the augmented control problem.

III. TIME CONSISTENCY AND CONCEPTUAL RISK-TO-GO

We start this section by demonstrating some time-inconsistency behaviors in risk-sensitive optimal control problems using several counter-examples.

A. Time Inconsistent Planning Leads to Irrational Behaviors

The most common strategy to model risk awareness in MDPs is to consider a static risk metric (i.e., a metric assessing risk from the perspective of a *single* point in time) applied to the entire stream of future costs. Typical examples include variance-constrained MDPs [17], [18], [19], or problems with probability constraints [17], [20], [21], [22], which are popular in the robotics community (in these problems risk is assessed only from the perspective of the initial stage). However, since static risk metrics do not involve a reassessment of risk at subsequent decision stages they generally lead to *irrational* behaviors. For example, a UV can seek to incur losses (i.e., dangerous maneuvers) or can deem as dangerous states that are indeed favorable under any realization of the underlying uncertainty.

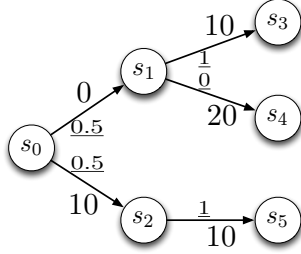
In this subsection, we will illustrate some irregular behaviors in risk sensitive multi-period planning by two examples.

Example 1: Variance-constrained planning — Given an MDP with initial state $x_0 \in S$ and time horizon $N \geq 1$, solve

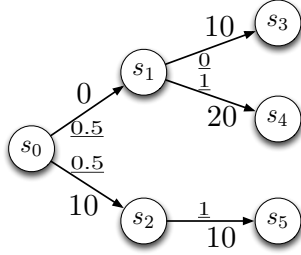
$$\begin{aligned} \min_{\pi} \quad & \mathbb{E} \left[\sum_{k=0}^{N-1} c(x_k, u_k) + c_N(x_N) \right] \\ \text{subject to} \quad & \text{var} \left(\sum_{k=0}^{N-1} d(x_k, u_k) + d_N(x_N) \right) \leq r_0, \end{aligned}$$

where $r_0 \in \mathbb{R}$ is a user-provided risk threshold.

Consider the example in Figure 1. When the risk threshold r_0 is below 25, policy π_1 is infeasible and the optimal policy is π_2 . According to policy π_2 , if the decision maker does not incur a cost in the first stage it *seeks to incur losses* in subsequent stages to keep the variance small. This can be seen as a consequence of the fact that Bellman's principle of optimality does not hold for this class of problems.



(a) Stage-wise constraint and objective function costs and transition probabilities for policy π_1 .



(b) Stage-wise constraint and objective function costs and transition probabilities for policy π_2 .

Fig. 1. Limitations of mean-variance optimization. Underlined numbers along the edges represent transition probabilities; non-underlined numbers represent stage-wise constraint and objective function costs (that are equal for this example). Terminal constraint costs are zero. Under policy π_1 , the costs per stage are given by $d(s_0, u_0) = 0.5 \cdot 0 + 0.5 \cdot 10 = 5$, $d(s_1, u_1) = 10$, and $d(s_2, u_1) = 10$; under policy π_2 , the costs per stage are given by $d(s_0, u_0) = 5$, $d(s_1, u_1) = 20$, and $d(s_2, u_1) = 10$. One can verify that for policy π_1 one has $\text{var}\left(\sum_{k=0}^{N-1} d(x_k, u_k) + d(x_N)\right) = 25$, while for policy π_2 one has $\text{var}\left(\sum_{k=0}^{N-1} d(x_k, u_k) + d(x_N)\right) = 0$. Then, if the risk threshold is less than 25, the decision-maker would choose policy π_2 and would seek to incur losses in order to keep the variance small enough.

As a second example, we consider MDPs with average value at risk (AVaR) constraints, which are closely related to chance (i.e., probability) constraints and are enjoying a growing popularity, especially in the finance industry [23], due to favorable computational aspects, such as convexity. The average value at risk for a random variable X at confidence level α is defined as [24]:

$$\text{AVaR}_\alpha(X) := \frac{1}{1-\alpha} \int_\alpha^1 \text{VaR}_\tau(X) d\tau,$$

where $\text{VaR}_\alpha(X)$ is simply the α -quantile of random variable X , i.e.,

$$\text{VaR}_\alpha(X) = \min\{x \mid \mathbb{P}[X \leq x] \geq \alpha\}.$$

Intuitively, the AVaR_α is the expectation of X in the conditional distribution of its upper α -tail. For this reason,

it can be interpreted as a metric of “how bad is bad.” The risk metric AVaR_α is closely related to chance constraints, since the constraint $\text{VaR}_\alpha(X) \leq 0$ corresponds to the chance constraint $\mathbb{P}[X \leq 0] \geq \alpha$ [24].

Example 2: AVaR-constrained planning — Given an MDP with initial state $x_0 \in S$ and time horizon $N \geq 1$, solve

$$\begin{aligned} \min_{\pi} \quad & \mathbb{E} \left[\sum_{k=0}^{N-1} c(x_k, u_k) + c_N(x_N) \right] \\ \text{subject to} \quad & \text{AVaR}_\alpha \left(\sum_{k=0}^{N-1} d(x_k, u_k) + d_N(x_N) \right) \leq r_0, \end{aligned}$$

where $r_0 \in \mathbb{R}$ is a user-provided risk threshold.

Let us interpret the constraint costs d as acceptable if negative and unacceptable otherwise. Accordingly consider the example in Figure 2 (based upon [25]), with threshold $r_0 = 0$ and confidence level $1/3$. One can show that the problem (consisting of a single policy) is infeasible, since at the first stage AVaR is positive. On the other hand, the constraint costs are acceptable in every state of the world from the perspective of the subsequent stage. In other words, the decision-maker would deem infeasible a problem that, at the second stage, appears feasible under any possible realization of the uncertainties.

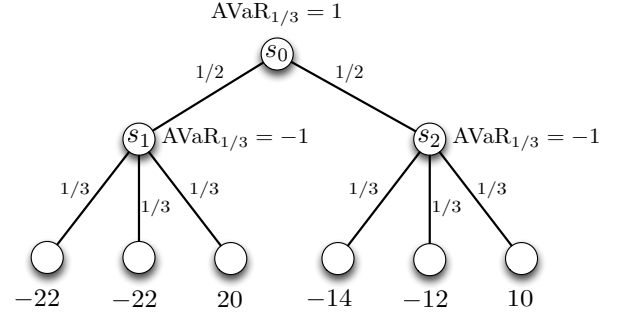


Fig. 2. Limitations of AVaR-constrained optimization. The numbers along the edges represent transition probabilities, while the numbers below the terminal nodes represent the terminal constraint costs (the other constraint costs are zero). The problem involves a single control policy (hence there is a unique transition graph). The constraint cost appears acceptable in states s_1 and s_2 , but unacceptable from the perspective of the first stage in state s_0 .

It is important to note that there is nothing special about these examples, which indeed capture a range of widely accepted criteria to pose risk constraints. Similar paradoxical results could be obtained with other risk metrics. Henceforth, we will collectively refer to the aforementioned irrational behaviors as “time-inconsistent” policies, since they reflect an inconsistent risk assessment over time.

B. Time Consistency

From Theorem II.4 and Theorem II.5 we can find a sequence of history dependent optimal control policies by Bellman iteration. In this section, we want to show that the risk constrained SOC in Problem OPT is time consistent. Most analysis in literature (c.f. Chapter 1 of [26], [27] for more details), restricts the analysis of time consistency to problems with Markovian policies. Since Problem OPT solves for a sequence of history dependent policy, it is unclear how we can analyze time consistency directly by

truncating the sub-histories. For this purpose, we define the space of augmented state feedback policies and the space of risk-to-go updates:

$$\Pi_k^m := \left\{ \{\pi_k, \dots, \pi_{N-1}\} : \begin{array}{l} \pi_j : S \times \Phi_j(S) \rightarrow U(S), \\ j \in \{k, \dots, N-1\} \end{array} \right\}$$

$$\mathfrak{R}_k := \left\{ \{\mathcal{R}_{k+1}, \dots, \mathcal{R}_N\} : \begin{array}{l} \mathcal{R}_{j+1}(x_j, r_j)(x_{j+1}) = r_{j+1} - r_j, \\ r_{j+1} \in \Phi_{j+1}(x_{j+1}), \\ j \in \{k, \dots, N-1\} \end{array} \right\}$$

and the following optimization problem:

Optimization problem \mathcal{OPT}^M — Given an initial state $x_0 \in S$, a time horizon $N \geq 1$, and a risk threshold $r_0 \in \mathbb{R}$, solve

$$\begin{aligned} & \min_{\pi \in \Pi_0^m, \mathcal{R} \in \mathfrak{R}_0} J_N^\pi(x_0) \\ & \text{subject to} \quad \rho_j(r_{j+1} + d_j(x_j, \pi_j(x_j, r_j))) = r_j, \\ & \quad r_{j+1} = \mathcal{R}_{j+1}(x_j, r_j)(x_{j+1}) + r_j \\ & \quad 0 \leq r_N, \quad j \in \{0, \dots, N-1\}. \end{aligned}$$

The k -th subproblem of Problem \mathcal{OPT}^M is simply defined by replacing $x_0 \in S$ and $r_0 \in \Phi_0(x_0)$ by $x_k \in S$ and $r_k \in \Phi_k(x_k)$, i.e.,

- If $k < N$ and $r_k \in \Phi_k(x_k)$:

$$\begin{aligned} & V_k^M(x_k, r_k) = \\ & \min_{\pi \in \Pi_k^m, \mathcal{R} \in \mathfrak{R}_k} J_N^\pi(x_k) \\ & \text{subject to} \quad \rho_j(r_{j+1} + d_j(x_j, \pi_j(x_j, r_j))) = 0, \\ & \quad r_{j+1} = \mathcal{R}_{j+1}(x_j, r_j)(x_{j+1}) + r_j \\ & \quad 0 \leq r_N, \quad j \in \{k, \dots, N-1\}. \end{aligned}$$

- If $k \leq N$ and $r_k \notin \Phi_k(x_k)$:

$$V_k^M(x_k, r_k) = \infty.$$

- When $k = N$ and $r_N \in \Phi_N(x_N) = [0, \infty]$:

$$V_N^M(x_N, r_N) = 0.$$

We are now in the position of defining the notion of time consistency for Problem \mathcal{OPT}^M .

Definition III.1. (Time Consistency of risk constrained SOC) For any given initial state $x_0 \in S$, and risk threshold $r_0 \in \Phi_0(x_0)$, define $\pi^* = \{\pi_0^*, \dots, \pi_{N-1}^*\} \in \Pi_0^M$ as a sequence of optimal policy and $r^* = \{r_1^*, \dots, r_N^*\}$ as a sequence of risk-to-go for Problem \mathcal{OPT}^M . Problem \mathcal{OPT}^M is a time consistent SOC problem, if at any $k \in \{0, \dots, N-1\}$, the k -subsequence of π^* : $\pi_{k,N-1}^* := \{\pi_k^*, \dots, \pi_{N-1}^*\}$ and r^* : $r_{k+1,N}^* := \{r_{k+1}^*, \dots, r_N^*\}$ are sequences of optimal solution to the k -tail subproblem of Problem \mathcal{OPT}^M at initial state $x_k \in S$, and risk threshold $r_k \in \Phi_k(x_k)$. In this case, π^* is a time consistent optimal policy, and \mathcal{R}^* is a time consistent risk-to-go.

Before getting into the main result, we want to justify the equivalence between Problem \mathcal{OPT}^M and \mathcal{OPT} . First, we have the following technical lemma, showing that without loss of optimality, the inequality constraint in $F(x_k, r_k)$ can be replaced by an equality.

Lemma III.2. For any value function $V_{k+1} : S \times \Phi_{k+1}(S) \rightarrow \mathbb{R}$, $\forall k \in \{0, \dots, N-1\}$ for Problem \mathcal{OPT} , the following equality holds:

$$V_k(x_k, r_k) = T_k[V_{k+1}](x_k, r_k) = \bar{T}_k[V_{k+1}](x_k, r_k), \quad \forall x_k, r_k. \quad (6)$$

where

$$\bar{T}_k[V](x_k, r_k) := \inf_{(u, r') \in \bar{F}_k(x_k, r_k)} \left\{ c(x_k, u) + \sum_{x_{k+1} \in S} Q(x_{k+1}|x_k, u) V(x_{k+1}, r'(x_{k+1})) \right\},$$

and \bar{F}_k is the set of control/threshold functions:

$$\bar{F}_k(x_k, r_k) := \left\{ (u, r') \mid u \in U(x_k), r'(x') \in \Phi_{k+1}(x') \text{ for all } x' \in S, r'(x') = r_k + L(x'), \text{ and } \rho_k(L) + d(x_k, u) = 0 \right\}.$$

Proof. First it is obvious that $\bar{F}_k(x_k, r_k) \subseteq F_k(x_k, r_k)$ and $T_k[V_{k+1}](x_k, r_k) \leq \bar{T}_k[V_{k+1}](x_k, r_k)$. Now, suppose $F_k(x_k, r_k) \neq \emptyset$ and there exists an risk-to-go $r'^*(x_{k+1})$ such that

$$d(x_k, u^*) + \rho_k(r'^*(x_{k+1})) < r_k$$

where u^* is the optimal control input solved from the Bellman's equation in (5). Recall the definition of the value function of Problem \mathcal{OPT} , $V_{k+1}(x_{k+1}, r_{k+1}) = \min\{J_N^\pi(x_{k+1}) : \pi \in \Pi_{k+1}, R_N^\pi(x_{k+1}) \leq r_{k+1}\}$ when $r_{k+1} \in \Phi_{k+1}(x_{k+1})$. It can be easily seen that $V_{k+1}(x_{k+1}, r_{k+1})$ is a non-increasing function with respect to r_{k+1} . Furthermore, since ρ_k is Lipschitz and $r'^*(x_{k+1})$ is a discrete state, continuous magnitude random variable, we can always find a nonnegative discrete state, continuous magnitude, bounded random variable $\epsilon(x_{k+1})$ such that

$$d(x_k, u^*) + \rho_k(r'^*(x_{k+1}) + \epsilon(x_{k+1})) = r_k.$$

Furthermore, we know that $r'^*(x_{k+1}) + \epsilon(x_{k+1}) \geq r'^*(x_{k+1})$ surely, which implies $r'^*(x_{k+1}) + \epsilon(x_{k+1}) \geq \underline{R}_N(x_{k+1})$. On the other hand, since $r_k, \epsilon(x_{k+1})$ and $r'^*(x_{k+1})$ are all bounded quantities, this implies $r'^*(x_{k+1}) + \epsilon(x_{k+1}) < \infty$ surely. By writing $\tilde{r}'^*(x_{k+1}) = r'^*(x_{k+1}) + \epsilon(x_{k+1}) \geq r'^*(x_{k+1})$, one obtains $V_{k+1}(x_{k+1}, r'^*(x_{k+1})) \geq V_{k+1}(x_{k+1}, \tilde{r}'^*(x_{k+1}))$ surely. Thus, we can use $(u^*, \tilde{r}'^*) \in \bar{F}_k(x_k, r_k)$ as a minimizer for the Bellman's equation in (5). To summarize, whenever $F_k(x_k, r_k) \neq \emptyset$, there always exists an optimal control input u^* and risk-to-go $r'^*(x_{k+1})$ for Bellman's equation in (5) such that $d(x_k, u^*) + \rho_k(r'^*(x_{k+1})) = r_k$. By letting $L(x') = r'^*(x') - r_k, \forall x' \in S$, we have just showed that there is no loss of optimality to consider the operator $\bar{T}_k[V_{k+1}](x_k, r_k)$ instead of the Bellman's equation in (5). \square

Essentially, we have the following result. The proof of this theorem is analogous to the proof of Theorem II.4 and is omitted here for the interest of brevity. Details of the proof can be found in the Appendix.

Theorem III.3. The value function of Problem \mathcal{OPT}^M is identical to the value function of Problem \mathcal{OPT} . Furthermore, Problem \mathcal{OPT}^M is a time consistent SOC problem.

Assume the infimum in expression (6) is attained. Let $r_0^* = r_0$. Based on Definition III.1, any optimal control policies and optimal risk-to-go found from the sequence of Bellman's equation, $\forall k \in \{0, \dots, N-1\}$:

$$(\pi_k^*(x_k, r_k^*), r_{k+1}^*) \in \operatorname{argmin}_{(u, r') \in \bar{F}_k(x_k, r_k^*)} \left\{ c(x_k, u) + \sum_{x_{k+1} \in S} Q(x_{k+1}|x_k, u) V_{k+1}^M(x_{k+1}, r'(x_{k+1})) \right\}$$

are time consistent.

Remark III.4. Notice that the risk-to-go satisfies the following equation:

$$\rho_k(r_{k+1}^*) + d_k(x_k, \pi_k^*(x_k, r_k^*)) = r_k^*, \quad \forall k \in \{0, \dots, N-2\}.$$

where $r_0^* = r_0$. Define

$$\bar{M}_{k+1} = r_{k+1}^* + \sum_{j=0}^k d_j(x_j, \pi_j^*(x_j, r_j^*))$$

We can show that M_{N-1} satisfies the following risk sensitive Martingale property: $\rho_k(M_{k+1}) = M_k$.

Next, we have the following corollary depicting the closed form solution policy of Problem OPT. The proof is identical to the proof of Theorem II.5 and will be omitted for brevity.

Corollary III.5. Let $\pi \in \Pi$ be a policy recursively defined by the solution of Bellman's equation in (6):

$$\pi_k(h_k) = u^*(x_k, r_k^*), \quad \text{with } r_k^* = r_{k-1}^* + L^*(x_{k-1}, r_{k-1}^*)(x_k),$$

when $k \in \{1, \dots, N-1\}$, and

$$\pi_0(x_0) = u^*(x_0, r_0^*),$$

where $r_0^* = r_0 \in \Phi_0(x_0)$. Then, $\pi = \{\pi_0, \dots, \pi_{N-1}\}$ is an optimal policy for Problem OPT with initial state x_0 and threshold r_0 . Furthermore, the k -th subsequence of π , i.e., $\{\pi_k, \dots, \pi_{N-1}\}$ is also an optimal history dependent policy to the tail subproblem of Problem OPT.

Similar to problem OPT, this corollary concludes that the optimal policy of problem OPT^M can also be constructed using both state update and the risk-to-go. The resultant policy is thus history dependent. However it is still unclear how one can obtain the analytical formula of the risk-to-go from merely solving the Bellman iteration in (6).

C. Analytical Formula for the Risk-to-go Update

Based on the time consistency analysis in previous sections, we aim to derive the analytical update formula for the risk-to-go r_k . Before getting to the main result, we need the following technical lemmas.

Lemma III.6. Suppose $r_k = R_N^\pi(x_k) - R_N^\pi(x_{k-1}) + r_{k-1}$ for $k \in \{1, \dots, N-1\}$ and $\pi \in \Pi$ is any admissible history dependent control policy. Then, if $r_0 \in \Phi_0(x_0)$ such that $R_N^\pi(x_0) \leq r_0$, it implies $r_k \in \Phi_k(x_k)$ for all $k \in \{0, \dots, N-1\}$.

Proof. First, we characterize the lower bound of r_k . By definition, $r_{j+1} - r_j = R_N^\pi(x_{j+1}) - R_N^\pi(x_j)$, for $j \in$

$\{0, \dots, N-2\}$. By summing over $j = 0$ to $k-1$, for any fixed $k \in \{0, \dots, N-1\}$, we get

$$\begin{aligned} r_k - r_0 &= R_N^\pi(x_k) - R_N^\pi(x_0) \\ \implies r_k - R_N^\pi(x_k) &= r_0 - R_N^\pi(x_0) \end{aligned}$$

Thus, one obtains

$$r_k - \underline{R}_N(x_k) \geq r_k - R_N^\pi(x_k) = r_0 - R_N^\pi(x_0) \geq 0.$$

The last inequality is due to the fact that $\pi \in \Pi$ is a feasible control policy. Next, we characterize the upper bound for r_k . For $r_k = R_N^\pi(x_k) - R_N^\pi(x_{k-1}) + r_{k-1}$, by monotonicity and translation invariance of multi period risk measures, one can easily show that for any $k \in \{0, \dots, N-1\}$,

$$R_N^\pi(x_k) \leq (N-k)\rho_{\max}, \quad R_N^\pi(x_k) \geq (N-k)\rho_{\min}.$$

Therefore, the above expressions imply

$$r_k - r_{k-1} \leq (N-k)(\rho_{\max} - \rho_{\min}) - \rho_{\min}.$$

By a telescopic sum, and since $r_0 \in [\underline{R}_N(x_0), \infty)$, one obtains

$$\begin{aligned} r_k &\leq \sum_{j=1}^k (N-j)(\rho_{\max} - \rho_{\min}) - k\rho_{\min} + r_0 \\ &= r_0 + k \left(\left(N - \frac{k+1}{2} \right) (\rho_{\max} - \rho_{\min}) - \rho_{\min} \right) < \infty \end{aligned}$$

Thus, combining the result of lower bound for r_k , we get, $r_k \in [\underline{R}_N(x_k), \infty)$, which means $r_k \in \Phi_k(x_k)$ for $k \in \{0, \dots, N-1\}$. \square

We are now in the position of deriving the main result of this paper. The following theorem provides an analytical update formula for the risk-to-go r_k .

Theorem III.7. (Conceptual Risk-to-go) Let $\pi^* \in \Pi$ be an optimal policy for Problem OPT. The following risk-to-go

$$\begin{aligned} \tilde{r}_k &= r_0, \\ \tilde{r}_{k+1} &= \tilde{r}_k + R_N^{\pi^*}(x_{k+1}) - R_N^{\pi^*}(x_k), \quad \forall k \in \{0, \dots, N-1\}, \end{aligned} \tag{7}$$

and the augmented state-feedback control policy $\tilde{\pi} = \{\tilde{\pi}_0, \dots, \tilde{\pi}_{N-1}\}$, where $\tilde{\pi}(x_k, \tilde{r}_k) = \pi^*(x_k)$ form time consistent solution to Problem OPT^M.

Proof. Since $R_N^{\pi^*}(x_k)$ and $R_N^{\pi^*}(x_{k+1})$ are formed by compounding Markov risk measures, we can easily see from Definition II.3 that $R_N^{\pi^*}(x_k)$ and $R_N^{\pi^*}(x_{k+1})$ are functions of x_k and x_{k+1} . Also, define

$$\tilde{L}_k(x_{k+1}) = R_N^{\pi^*}(x_{k+1}) - R_N^{\pi^*}(x_k),$$

for any $k \in \{0, \dots, N-1\}$. From Lemma III.6, one obtains $\tilde{r}_{k+1} = \tilde{r}_k + \tilde{L}_k(x_{k+1}) \in \Phi_{k+1}(x_{k+1})$, $\forall k \in \{0, \dots, N-2\}$. Also, as there is no terminal cost, by a telescopic sum,

$$\tilde{r}_N = \tilde{r}_0 + \sum_{k=0}^{N-1} \tilde{L}_k(x_{k+1}) = -R_N^{\pi^*}(x_0) + \tilde{r}_0.$$

As π^* is an optimal history dependent control policy for Problem OPT, one obtains $R_N^{\pi^*}(x_0) \leq \tilde{r}_0$ and $\tilde{r}_N \geq 0$. Furthermore, by the time consistent property and translational invariance of risk measures, we have that

$$\begin{aligned} \rho_k(R_N^{\pi^*}(x_{k+1})) + c(x_k, \pi_k^*(x_k)) &= R_N^{\pi^*}(x_k), \quad \forall k, \\ \iff \rho_k(R_N^{\pi^*}(x_{k+1})) + c(x_k, \tilde{\pi}_k(x_k, \tilde{r}_k)) &= R_N^{\pi^*}(x_k), \quad \forall k, \end{aligned}$$

Thus, the optimal policy $\tilde{\pi}$ and the risk-to-go sequence $\tilde{r} = \{\tilde{r}_1, \dots, \tilde{r}_N\}$ are feasible to Problem \mathcal{OPT}^M . Then, $\forall k \in \{0, \dots, N-1\}$,

$$V_0(x_0, r_0) = J_N^{\pi^*}(x_0) = J_N^{\tilde{\pi}}(x_0) \geq V_0^M(x_0, r_0).$$

At the same time, equation (8) implies that $V_0^M(x_0, r_0) \geq V_0(x_0, r_0)$. Thus both arguments imply $\tilde{\pi}$ and $\tilde{r} = \{\tilde{r}_1, \dots, \tilde{r}_N\}$ form solution to Problem \mathcal{OPT}^M . Time consistency then follows directly from Definition III.1. \square

From this theorem, we conclude that the analytical formula of the risk-to-go can be written as a Martingale difference of the constraint cost function. This property is crucial to understand how risk evaluation is updated at each step in order to make time consistent decisions and to derive large scale risk constrained decision making algorithms.

IV. THE SQUANDER OR SAVE EXAMPLE (C.F. [2])

We consider the simple case of risk constrained SOC problem where risk-neutral costs and risk neutral constraints are considered. Given an MDP with initial state $x_0 \in S_0$, risk threshold $r_0 \in \Phi_0(x_0)$ and time horizon $N \geq 1$, solve

$$\begin{aligned} \min_{\pi} \quad & \mathbb{E} \left[\sum_{k=0}^{N-1} c(x_k, u_k) \right] \\ \text{subject to} \quad & \mathbb{E} \left[\sum_{k=0}^{N-1} d(x_k, u_k) \right] \leq r_0, \end{aligned}$$

Consider the example in Figure 3 and Figure 4 with $N = 2$, and all terminal costs equal to zero (they are not drawn in the figures). Also, let S_0 be the initial state space,

$$\begin{aligned} S_1 &= \{\text{"win a lottery"}, \text{"lose a lottery"}\}, \\ U(\text{"win a lottery"}) &= \{\text{"squander"}, \text{"save"}\}, \\ U(\text{"lose a lottery"}) &= \{\text{"squander"}, \text{"save"}\} \end{aligned}$$

be the state and closed loop action space at stage 1, and

$$S_2 = \left\{ \begin{array}{l} \text{"win a lottery and squander"}, \\ \text{"win a lottery and save"}, \\ \text{"lose a lottery and squander"}, \\ \text{"lose a lottery and save"} \end{array} \right\},$$

be the state space at stage 2. Suppose there are no actions ($U(S_0) = \{\emptyset\}$), no cost ($c(x_0, u_0) = 0, \forall (x_0, u_0) \in S_0 \times U(S_0)$) and no constraint cost ($d(x_0, u_0) = 0, \forall (x_0, u_0) \in S_0 \times U(S_0)$) for winning/losing a lottery at stage 0. On the other hand, the stage-wise cost in stage 1 is as follows:

$$\begin{aligned} c(\text{"win a lottery"}, u_1) &= \begin{cases} -50 & \text{if } u_1 = \text{"squander"} \\ -30 & \text{if } u_1 = \text{"save"} \end{cases} \\ c(\text{"lose a lottery"}, u_1) &= \begin{cases} -20 & \text{if } u_1 = \text{"squander"} \\ -10 & \text{if } u_1 = \text{"save"} \end{cases} \end{aligned}$$

and the constraint stage-wise cost is as follows:

$$\begin{aligned} d(\text{"win a lottery"}, u_1) &= \begin{cases} 1 & \text{if } u_1 = \text{"squander"} \\ 0.05 & \text{if } u_1 = \text{"save"} \end{cases} \\ d(\text{"lose a lottery"}, u_1) &= \begin{cases} 0.4 & \text{if } u_1 = \text{"squander"} \\ 0.2 & \text{if } u_1 = \text{"save"} \end{cases} \end{aligned}$$

Suppose at time 0, one has probability ϵ of winning a lottery and probability $1 - \epsilon$ of losing a lottery where $0 < \epsilon < 1$. If one wins the lottery, at stage 1, one can

choose to squander (action 1) or to save (action 2). If one loses the lottery, at stage 1 one can choose to squander (action 1) or to save (action 2). In this example, the stage-wise cost represents a level of satisfaction in spending and it is inversely proportional to the money spent. On the other hand, the stage-wise constraint cost is the probability of going bankrupt, which is directly proportional to the money spent.

Consider $\epsilon = 0.1$ and the risk threshold $r_0 = 0.3$. That is, the probability of winning lottery is 0.1 and one wants to limit the probability of bankruptcy to be under 0.3. Similar to Haviv's argument in [2], if we keep the risk threshold constant at 0.3 for all subsequent stages, the optimal policy decided at stage 0 is not to squander if one loses the lottery, and to squander if one wins the lottery. However, the optimal policy decided in stage 1 is to save if one loses or wins the lottery. This implies that the optimal control policies decided at stage 0 is time-inconsistent.

On the other hand, suppose one finds the optimal control policies by solving the Bellman's equation in (5). Then, from the value function for Problem \mathcal{OPT} , one obtains $r_2 = 0$ and $V(x_2, r_2) = 0$. Now let s_{11} = "win a lottery" and s_{12} = "lose a lottery". At stage 1, the value function is as follows:

$$\begin{aligned} V_1(s_{11}, r_1(s_{11})) &= \begin{cases} -50 & \text{if } r_1(s_{11}) \geq 1 \\ -30 & \text{if } 0.05 \leq r_1(s_{11}) < 1 \\ \infty & \text{otherwise} \end{cases} \\ V_1(s_{12}, r_1(s_{12})) &= \begin{cases} -20 & \text{if } r_1(s_{12}) \geq 0.4 \\ -10 & \text{if } 0.2 \leq r_1(s_{12}) < 0.4 \\ \infty & \text{otherwise} \end{cases} \end{aligned}$$

and the optimal policy is as follows:

$$\begin{aligned} \pi_1^*(s_{11}, r_1(s_{11})) &= \begin{cases} \text{"squander"} & \text{if } r_1(s_{11}) \geq 1 \\ \text{"save"} & \text{if } 0.05 \leq r_1(s_{11}) < 1 \\ \emptyset & \text{otherwise} \end{cases} \\ \pi_1^*(s_{12}, r_1(s_{12})) &= \begin{cases} \text{"squander"} & \text{if } r_1(s_{12}) \geq 0.4 \\ \text{"save"} & \text{if } 0.2 \leq r_1(s_{12}) < 0.4 \\ \emptyset & \text{otherwise} \end{cases} \end{aligned}$$

At stage 0, with $r_0 = 0.3$, the value function is:

$$\begin{aligned} V_0(s_0, r_0) &= \min_{r_1(\cdot) \in \mathbb{R}^2} \left[0.1 \times \begin{cases} -50 & \text{if } r_1(s_{11}) \geq 1 \\ -30 & \text{if } 0.05 \leq r_1(s_{11}) < 1 \\ \infty & \text{if } r_1(s_{11}) < 0.05 \end{cases} \right. \\ &\quad \left. + 0.9 \times \begin{cases} -20 & \text{if } r_1(s_{12}) \geq 0.4 \\ -10 & \text{if } 0.2 \leq r_1(s_{12}) < 0.4 \\ \infty & \text{if } r_1(s_{12}) < 0.2 \end{cases} \right] \\ \text{subject to} \quad & r_1(s_{11}) \in [0.05, \infty], \quad r_1(s_{12}) \in [0.2, \infty], \\ & 0.1 \times r_1(s_{11}) + 0.9 \times r_1(s_{12}) \leq 0.3. \end{aligned}$$

This implies that the value function is

$$V_0(s_0, r_0) = -10(0.9) - 30(0.1) = -12.$$

Based on the risk-to-go update in equation (6), we have that

$$\begin{aligned} r^*(s_0, r_0)(s_{11}) &= 0.3 + 0.05 - (0.1(0.05) + 0.9(0.2)) = 0.165, \\ r^*(s_0, r_0)(s_{12}) &= 0.3 + 0.2 - (0.1(0.05) + 0.9(0.2)) = 0.315. \end{aligned}$$

Thus, the optimal history dependent policy at stage 0 is

$$\begin{aligned} \pi_0^*(h_0) &= \emptyset, \\ \pi_1^*(h_1) &= \begin{cases} \text{"save"} & \text{if } x_1 = \text{"win a lottery"} \\ \text{"save"} & \text{if } x_1 = \text{"lose a lottery"} \end{cases}, \end{aligned}$$

i.e., this optimal policy is time consistent.

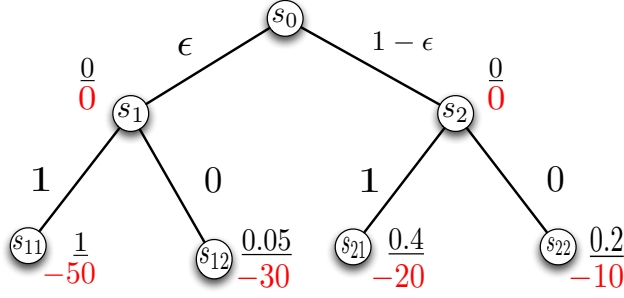


Fig. 3. Transition probability, cost function and constraint cost function under policy 1 (squander). The red numbers indicates the stage-wise cost. The underlined black numbers indicates the stage-wise constraint cost

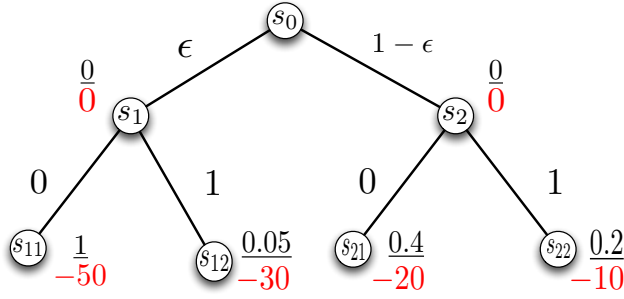


Fig. 4. Transition probability, cost function and constraint cost function under policy 2 (save). The red numbers indicates the stage-wise cost. The underlined black numbers indicates the stage-wise constraint cost.

V. CONCLUSION

In this paper we study time-consistency of risk constrained problem where the risk metric is time consistent. From the Bellman optimality condition in [1], we establish an analytical “risk-to-go” that results in a time consistent optimal policy. The effectiveness of the analytical solution is also justified by solving Haviv’s counter-example [2] in time inconsistent planning. Future work includes extending the above analysis to large scale risk constrained decision making (via the use of approximate dynamic programming) that provides time consistent policies.

REFERENCES

- [1] Y. Chow and M. Pavone. Stochastic optimal control with dynamic, time-consistent risk constraints. In *American Control Conference*, 2013.
- [2] M. Haviv. On constrained Markov decision processes. *Operations research letters*, 19(1):25–28, 1996.
- [3] P. Artzner, F. Delbaen, J. Eber, D. Heath, and H. Ku. Coherent multiperiod risk adjusted values and Bellman’s principle. *Annals of Operations Research*, 152(1):5–22, 2007.
- [4] B. Acciaio, H. Follmer, and I. Penner. *Dynamic convex risk measures*, chapter 1, pages 1–34. Springer-Verlag Berlin Heidelberg, 2011.
- [5] B. Roorda, J. M. Schumacher, and J. Engwerda. Coherent acceptability measures in multi-period models. *Mathematical Finance*, 15(4):589–612, 2005.
- [6] M. Hardy and J. Wirch. The iterated CTE: a dynamic risk measure. *North American Actuarial Journal*, 8(4):62–75, 2004.
- [7] F. Riedel. Dynamic coherent risk measures. Technical report, Stanford University, 2003.
- [8] K. Boda and J. A. Filar. Time consistent dynamic risk measures. *Mathematical Methods in Operations Research*, 63(1):169–186, 2006.

- [9] H. Follmer and A. Schied. Coherent and convex risk measures. In *Encyclopedia of Quantitative Finance*, number 355-363. John Wiley & Sons, 2010.
- [10] A. Shapiro. Minimax and risk averse multistage stochastic programming. *European Journal of Operational Research*, 219(3):719–726, 2012.
- [11] A. Ruszczyński. Risk averse dynamic programming for Markov decision process. *Journal of Mathematical Programming*, 125(2):235–261, 2010.
- [12] D. A. Iancu, M. Petrik, and D. Subramanian. Tight approximations of dynamic risk measures. 2013. Submitted, available at <http://arxiv.org/abs/1106.6102>.
- [13] T. Osogami and T. Morimura. Time-consistency of optimization problems. Technical report, IBM Japan, 2010.
- [14] K. Ross and R. Varadarajan. Markov decision processes with sample path constraints: the communicating case. *Operations Research*, 37(5):780–790, 1989.
- [15] L. Sennott. Another set of conditions for average optimality in Markov control processes. *Systems & control letters*, 24(2):147–151, 1995.
- [16] R. Strotz. Myopia and inconsistency in dynamic utility maximization. *The Review of Economic Studies*, pages 165–180, 1955.
- [17] A. Piunovskiy. Dynamic programming in constrained Markov decision process. *Control and Cybernetics*, 35(3):646–660, 2006.
- [18] M. Sniedovich. A variance-constrained reservoir control problem. *Water Resources Research*, 16:271–274, 1980.
- [19] S. Mannor and J. N. Tsitsiklis. Mean-variance optimization in Markov decision processes. In *International Conference on Machine Learning*, 2011.
- [20] M. Ono and B. Williams. Iterative risk allocation: A new approach to robust model predictive control with a joint chance constraint. In *Decision and Control, 2008. CDC 2008. 47th IEEE Conference on*, pages 3427–3432. IEEE, 2008.
- [21] L. Blackmore, M. Ono, A. Bektassov, and B. Williams. A probabilistic particle-control approximation of chance-constrained stochastic predictive control. *Robotics, IEEE Transactions on*, 26(3):502–517, 2010.
- [22] Lars Blackmore and Masahiro Ono. Convex chance constrained predictive control without sampling. In *Proceedings of the AIAA Guidance, Navigation and Control Conference*, pages 7–21, 2009.
- [23] V. Borkar and R. Jain. Risk-constrained Markov decision processes. In *Decision and Control (CDC), 2010 49th IEEE Conference on*, pages 2664–2669. IEEE, 2010.
- [24] R. Rockafellar and S. Uryasev. Conditional value-at-risk for general loss distributions. *Journal of Banking & Finance*, 26(7):1443–1471, 2002.
- [25] P. Artzner, F. Delbaen, J. Eber, and D. Heath. Coherent measures of risk. *Mathematical finance*, 9(3):203–228, 1998.
- [26] D. Bertsekas. *Dynamic programming and optimal control*. Athena Scientific, 2005.
- [27] E. Chong, S. Miller, and J. Adaska. On Bellman’s principle with inequality constraints. *Operations Research Letters*, 40(2):108–113, 2012.

APPENDIX

A. Proof of Theorem III.3

The proof follows from the Bellman's equation with risk constraints. First, we want to prove that

$$V_k^M(x_k, r_k) = V_k(x_k, r_k) \quad (8)$$

for any $k \in \{0, \dots, N-1\}$, $x_k \in S$ and $r_k \in \Phi_k(x_k)$ and $k \in \{0, \dots, N-1\}$. At $k = N-1$, by definition, $V_{N-1}^M(x_{N-1}, r_{N-1}) = \bar{T}_{N-1}[V_N^M](x_{N-1}, r_{N-1}) = \bar{T}_{N-1}[V_N](x_{N-1}, r_{N-1})$. Since $V_N(x_N, r_N) = V_N^M(x_N, r_N) = 0$ for any $r_N \in [0, \infty]$, by Lemma III.2, one obtains $V_{N-1}^M(x_{N-1}, r_{N-1}) = \bar{T}_{N-1}[V_N](x_{N-1}, r_{N-1}) = T_{N-1}[V_N](x_{N-1}, r_{N-1}) = V_{N-1}(x_{N-1}, r_{N-1})$. Thus, equation (8) holds for $k = N-1$. By inductive hypothesis, suppose for $k = j+1$,

$$V_{j+1}^M(x_{j+1}, r_{j+1}) = V_{j+1}(x_{j+1}, r_{j+1})$$

Then for $k = j$. Let $\pi^* \in \Pi_j^M$ be the optimal policy that yields the optimal cost $V_j^M(x_j, r_j)$ and r^* be the sequence of optimal risk threshold update functions. By applying the law of total expectation, we can write:

$$\begin{aligned} V_j^M(x_j, r_j) &= \mathbb{E} \left[\sum_{i=j}^{N-1} c_i(x_i, \pi_i^*(x_i, r_i)) \right] \\ &= \mathbb{E} [c_j(x_j, \pi_j^*(x_j, r_j)) + J_N^{\pi^*}(x_{j+1})]. \end{aligned}$$

Clearly, π^* is a feasible policy for the $j+1^{\text{th}}$ tail subproblem of Problem \mathcal{OPT}^M with $x_{j+1} \in S$ and $r_{j+1}^* \in \Phi_{j+1}(x_{j+1})$. Collecting the above results, we can write

$$\begin{aligned} V_j^M(x_j, r_j) &\geq \mathbb{E} [c_j(x_j, \pi_j^*(x_j, r_j)) + V_{j+1}^M(x_{j+1}, r_{j+1}^*)] \\ &\geq \bar{T}_j[V_{j+1}^M](x_j, r_j) = \bar{T}_j[V_{j+1}](x_j, r_j) \\ &= T_j[V_{j+1}](x_j, r_j) = V_j(x_j, r_j). \end{aligned}$$

The second inequality follows from the fact that

$$\rho_j(r_{j+1}^*) + d_j(x_j, \pi_j^*(x_j, r_j)) \leq r_j.$$

and the first equality follows from induction's assumption.

On the other hand, for given pair (x_j, r_j) , where $r_j \in \Phi_j(x_j)$, let $u^*(x_j, r_j)$ and $r'^*,*(x_j, r_j)(x_{j+1})$ be the minimizers in $\bar{T}_j[V_{j+1}^M](x_j, r_j)$. Construct a policy $\bar{\pi} \in \Pi_j^M$ as follows: $\bar{\pi}_j(x_j, r_j) = u^*(x_j, r_j)$ and $\bar{\pi}_i(x_i, r_i) = \pi_i^*(x_i, r_i)$ for $i \geq j+1$. Therefore, the policy $\bar{\pi} \in \Pi_j^M$ is a feasible policy for the $j+1^{\text{th}}$ tail subproblem of Problem \mathcal{OPT}^M with $r_j \in \Phi_j(x_j)$. Hence, one easily obtains:

$$\begin{aligned} V_j^M(x_j, r_j) &\leq J_N^{\bar{\pi}}(x_j) = \mathbb{E} [c_j(x_j, \bar{\pi}_j(x_j, r_j)) + J_N^{\pi^*}(x_{j+1})] \\ &= \bar{T}_j[V_{j+1}^M](x_j, r_j) = \bar{T}_j[V_{j+1}](x_j, r_j) = V_j(x_j, r_j). \end{aligned}$$

By combining both steps, the claim in equation (8) is proved by induction. Furthermore, by the equation (8) and theorem II.4, the following expressions hold:

$$\begin{aligned} V_k^M(x_k, r_k) &= V_k(x_k, r_k) = \bar{T}_k[V_{k+1}](x_k, r_k) \\ &= \bar{T}_k[V_{k+1}^M](x_k, r_k). \end{aligned} \quad (9)$$

By repeatedly applying this Bellman's equation, one obtains

$$V_0^M(x_0, r_0) = \bar{T}_0[\bar{T}_1[\dots[\bar{T}_{N-1}[V_N^M]\dots]](x_0, r_0), \quad (10)$$

Also, define the sequence of optimal policy: $\pi^* = \{\pi_0^*, \dots, \pi_{N-1}^*\}$ and sequence of risk-to-go: $r^* =$

$\{r_1^*, \dots, r_N^*\}$ found by solving the sequence of Bellman's recursions in equation (10). Now, consider the k subsequence: $\tilde{\pi}^* = \{\pi_k^*, \dots, \pi_{N-1}^*\} \in \Pi_k^M$ of π^* and the k subsequence: $\tilde{r}^* = \{r_{k+1}^*, \dots, r_N^*\}$. By definition, one obtains,

$$d(x_j, \pi_j^*(x_j, r_j)) + \rho_j(r_{j+1}^*) \leq r_j, \quad \forall N-1 \geq j \geq k.$$

and the Bellman's equation implies

$$J_N^{\tilde{\pi}^*}(x_k) = \bar{T}_k[\dots[\bar{T}_{N-1}[V_N^M]\dots]](x_k, r_k) = V_k^M(x_k, r_k).$$

This implies that $\tilde{\pi}^*$ is a sequence of optimal control policy of the k -tail subproblem of Problem \mathcal{OPT}^M , for any state $x_k \in S$, and any risk threshold $r_k \in \Phi_k(x_k)$.